A STOCHASTIC CONSTELLATION REPLENISHMENT PLANNER

F. Letterio[(1)], G. Vicario de Miguel[(1)], J. Martín[(1)], D. Navarro Reyes [(2)], A. Ayala [(3)]

(1) Deimos Space SLU, Madrid, Spain
(2) European Space Agency, Noordwijk, The Netherlands
(3) GMV, Madrid, Spain

ABSTRACT

Constellation replenishment will determine whether a planned constellation can deliver its objectives within budget over the long term. The key measure is service availability: the fraction of the time when there are sufficient functioning satellites in the correct constellation slots. An efficient replenishment strategy means minimising the number of replacement satellites and launches required. In the frame of ESA's European GNSS Evolution Programme, Deimos Space developed a stochastic simulator with the objective of computing and trading off different constellation replenishment plans.

The tool allows replenishment strategy candidates to be designed in terms of launchers, spacecraft types, launch scenarios and transfer to the final slot in the constellation. A comprehensive trade-off definition has been performed, and the most promising subset of candidates selected for further analysis using the Replenishment Planner. Different methods were used to characterise candidate strategies, from bibliography research to low-thrust transfer optimization.

The Replenishment Planner was then designed and implemented. A stochastic approach is used to calculate the probability of having a certain number of satellites in each constellation plane. The Replenishment Planner maintains the service availability above a user defined threshold by means of different launch strategies, using the selected strategy candidates as "building blocks". The tool considers a wide set of inputs, conditions and constraints:

- Satellites types, e.g. current satellites, and future new designs or user defined satellites.
- Launch strategies can include triggering corrective launches to replace failed satellites, preventive launches to anticipate the decrease of the service availability, preventive launches to avoid the loss of spare satellite capability, or a combination of the previous.
- Transfer strategies including direct injection, electric thrust transfer, and staggered separation to inject satellites into different orbital planes at different altitudes.
- Satellite and launcher reliability.
- Policies on maintaining spare capacity and satellite decommissioning.
- Minimum intervals between launches and time required to procure a launch.

Simulations covering 100 years have been run, considering the decay of the current satellites and their replacement, using different launchers, computing relocation of spare satellites and considering temporary outages. Study cases have been run to validate the novel approach implemented in the Constellation Replenishment Planner, yielding promising results.

The approach is not limited to GNSS, but applicable to any constellation, and we discuss extension to Earth Observation and communications mega-constellations.

1. INTRODUCTION

Any long-duration service provided by a constellation of satellites must take into account the need to replace individual satellites. The high costs of building and launching satellites must be traded off against the costs of having gaps in service availability. Given that the lifetime of an individual satellite is unknown, probabilistic methods are valuable in planning constellation replenishment.

This paper presents the Replenishment Planner, a highly flexible tool able to provide valuable statistical information about the replenishment needs of a constellation, assessing different launch scenarios and replenishment strategies in terms of the quality of the service, and the number of launches and satellites required. The simulator is also able to evaluate the sensitivity of the service availability and the replenishment plan to multiple parameters such as the reliability of the satellites, the minimum time allowed between successive launches, service outages, etc.

The study, design and development of the Replenishment Planner consisted of two phases. During the first phase, a set of launch scenario candidates were designed, characterized in terms of launchers, spacecraft types and transfer to the constellation final slot. A comprehensive trade-off definition was performed and a sub-set of candidates were selected as the most promising ones to be further analysed by means of the Replenishment Planner. Different methods were used to perform the launch scenario candidate characterization, spanning from bibliography research to low-thrust transfer optimization using an improved Edelbaum method which takes into account Earth eclipses.

In a second phase, the Replenishment Planner was thus designed and implemented, based on a stochastic approach in order to calculate the probability of having a certain number of satellites in each plane of a user-given constellation. These statistics combined with satellite and launcher reliability figures, an input provided by system studies, give the key figure to be maintained by the simulator: the service availability.

2. CONSTELLATION REPLENISHMENT

The objective of the Replenishment Planner is to evaluate the constellation status in terms of service availability and to provide a plan of launches in order to maintain this figure of merit above a user-defined threshold. This section specifies in more detail the problem that must be solved, and the constraints that apply.

2.1. Service Availability

The availability of the service provided by a constellation is the fraction of time when the service is available in the target areas. This is a figure of merit that depends on the geometry of the constellation and on the number of satellites available (i.e. satellites providing the service). Therefore, the service availability of the constellation can be defined as the sum over the number of satellites of the availability coefficients multiplied by the probability of having that number of available satellites. Computation of the probabilities is very time consuming, making a fully analytic approach unfeasible with the current state-of-the-art computational capabilities. Therefore, the Replenishment Planner uses a stochastic approach, simulating the evolution of a user-defined number of constellations to retrieve the statistical information.

2.2. Constellation Evolution

The constellations simulated by the Replenishment Planner are Walker constellations with a user-defined number of planes, and nominal and spare slots per plane. The satellites in nominal slots are providing services, whereas the ones in spare slots are used as backup satellites, being transferred or relocated to a nominal slot whenever a nominal satellite belonging to its plane dies (stops providing the main service of the constellation). The re-location time and the maximum number of satellites that can be simultaneously in-transfer are user-defined parameters.

2.3. Satellite Models

The satellite model collects information regarding the satellites to be used. Different satellite models can be set up by the user, for example to allow for current and next generation satellites.

The services provided by the satellite are related to reliability curves that can model different instruments or subsystems of the satellites. Thus, the death date of each service is a random variable defined by its reliability curve, giving the problem a statistical nature. The Replenishment Planner is able to model very complex spacecraft that provide multiple services; it is possible to define complex relations between services (e.g. services that depend on many instruments, some of them shared). Services can either be the “driving” or main service whose failure would trigger spacecraft retirement, or non-driving additional services.

The spacecraft model can have other characteristics that contribute to fidelity of the simulator. Each service can have outages that temporarily reduce the service availability that can be scheduled (happening at a known date) or unscheduled (defined just with the expected number of occurrences per year).

The user can also define other specific parameters, e.g. a decommissioning age for the satellites (i.e. the satellites are retired after a certain number of years) given as a normal distribution, or the time needed to build a satellite that model the on-ground availability of satellites.

2.4. Launch Scenarios

The launch scenario defines the number of satellites that can be carried by a launcher, their satellite model and the transfer time to the nominal orbit. In addition, each launch scenario considers launcher reliability and transfer reliability.

Launch scenarios are flexible enough to model direct injection, quasi-direct injections, low thrust orbit raising and staggered launches, i.e. launch of multiple satellites to different orbital planes with different transfer times.

2.5. Initial Constellation State

The constellation state before the start of the replenishment phase is given by a set of already deployed satellites at a certain date and the initial deployment plan, i.e. the schedule of launches already planned at that epoch. The already deployed satellites are characterized by their satellite model (see section 2.3), launch date, and position in the constellation (plane number and slot allocation, i.e. spare or nominal). Since these satellites are supposed to be alive at the beginning of the replenishment simulation, their reliability curves are automatically computed using Bayesian probability.

Each launch of the initial deployment plan is defined with its launch date, the target primary plane and the associated launch scenario (see section 2.4).

2.6. Long-term Changes in Satellite Design, Launcher Availability and Replenishment Constraints

The simulator has been designed in order to manage the replenishment problem split into different blocks of time, referred to as launch intervals. It is possible to define many of the simulation parameters specifically for each launch interval. As well as the replenishment strategies discussed below, the following parameters can be defined for each launch interval: the minimum time between two consecutive launches, the minimum time needed to perform a launch after it is requested, the maximum number of satellites that can be launched during the launch interval or the available satellites (the on-ground available stock) at the beginning of the launch interval. Being able to split the simulation into these separate intervals gives the tool the flexible to deal with very large time frame (more than 100 years of simulation), over which time the satellite models and therefore the constraints and conditions will evolve significantly.

3. TOOL IMPLEMENTATION

The Replenishment Planner has been designed as a highly flexible tool. It is capable of simulating the evolution of a generic constellation given its configuration (number of planes and number of nominal and spare slots per plane) and its initial state. The Replenishment Planner has been developed in C++11 and for Linux OS, but tested also in Windows. It is able to provide the outputs in NetCDF4 and ASCII formats.

The Replenishment Planner includes a highly customizable post-processing module compatible with Octave 4.2 and Matlab 2013b to load and plot the simulation results.

The simulator has been integrated in openSF [1], an ESA generic simulation framework, which provides capabilities to perform Monte Carlo, batch and parametric analysis as well as input edition and managing external tools for output post-processing (e.g. plot generation).

The following subsections discuss specific features of the implementation.

3.1. Predictor-Corrector Replenishment Planner

The tool can be used to model how a replenishment plan will perform when faced with a real constellation evolution. This is done using a Monte Carlo approach, simulating many possible evolutions of the constellation in parallel. Within each Monte Carlo case, the Replenishment Planner is run at certain times (user defined to consider procurement constraints) in advance of each launch, first to predict the launches that will be needed, and then to correct the plan as the constellation evolves. Comparing the actual deployment schedules which result across the whole set of Monte Carlo cases allows the user to check the sensitivity of a nominal replenishment plan to different parameters, and its robustness against the typical constraints of the launch scheduling process.

The scheduling process of the Predictor-Corrector Replenishment Planner is intended to reproduce the typical launch scheduling constraints. The launches are first predicted using the current state of the constellation and the Replenishment Planner, announcing them with enough anticipation. This predicted launch date is not the final one, since there are a maximum delay and advance times for the launch around the predicted date when it can be performed. Then, a given time before the first predicted launch date (called check time), the launch must be fixed and committed.

The computation of the final date of the launch, which must be within the interval defined by the maximum delay and advance times, is performed in the correction phase using the state of the constellation at the check time and the Replenishment Planner. The other launches computed during the correction phase will be taken as predicted ones in order to keep on with the scheduling process. If there is enough time (i.e. if the check time of the first predicted launch is after the arrival of the satellites), these dates will be refined with the state of the constellation after the arrival of the satellites from the committed launches.

Thus, the simulation starts with the prediction, by means of the Replenishment Planner, of the first set of launches. After that, in the correction phase, the first predicted launch is corrected and committed, for which the "real" constellation is evolved from the initial state until the user defined launch check time previous to the predicted launch date. Using this new current state of the "real" constellation, the Replenishment Planner is executed again retrieving the set of corrected launches. From that set, all the launches occurring within the user defined interval around the first predicted launch will be committed, inserting a launch at the end of the interval if it results empty. After the correction phase, the real constellation advances until the committed launches are taken into account by the real constellation (if possible); this state is then taken as the starting point of the following prediction phase. This prediction and correction sequence is repeated until the end of the simulation.

The outcome of a single Predictor-Corrector Replenishment Planner simulation is the evolution of a single constellation and the set of committed launches, which will differ from the nominal plan computed by the Replenishment Planner. Therefore, this use case offers relevant results when a Monte-Carlo analysis is performed, since it provides valuable statistical information regarding the robustness and sensitivity of the nominal plan.

The Predictor-Corrector Replenishment Planner has two different working modes. The sensitivity mode does not restrict the number of launches committed in the simulation, whereas the robustness mode forces the number of launches committed to be equal to the number of launches in the nominal plan. To see the difference between the two modes, consider a case where it happens that the real constellation performs worse than expected between the prediction and correction phases (i.e. some satellites die, compromising the quality of the service provided), needing multiple launches in the interval of the first predicted launch to recover the quality of the service. In the sensitivity mode these launches will be committed, but in the robustness mode only the first one would be committed, postponing the other launches until a date close enough to the second predicted launch.

Therefore, the sensitivity mode focuses on maintaining the service level, so the output of the Monte-Carlo analysis is the variability (sensitivity) of the nominal launch plan. On the other hand, the robustness mode is constrained by the number of launches, so the main outcome of the Monte-Carlo is the evolution of the service availability, which is more likely to violate the threshold level.

3.2. Parametric and Monte Carlo Analysis

The Replenishment Planner can be easily configured to perform parametric and Monte Carlo analysis, varying most of the input variables. Moreover the computation engine in charge of this feature is able to distribute the computational burden to a user-defined number of CPU cores, drastically increasing the simulator performances in terms of execution time.

These modes can be used to perform sensitivity analysis of the service availability and the replenishment plan to relevant inputs (e.g. satellites reliability, minimum time between launches, launch reliability, relocation time, decommissioning age, etc.). In addition, robustness analysis of a reference replenishment plan can be carried out setting it as the deployment plan and performing parametric or Monte Carlo simulations on it.

As stated in section 3.1, the Predictor-Corrector Replenishment Planner is especially useful when is combined with a Monte Carlo analysis, since it is able to provide relevant statistical information regarding the robustness and sensitivity of the nominal plan submitted to the typical launch scheduling constraints.

3.3. Outputs

The simulator is able to provide a variety of outputs such as the evolution of the service availability (with and without considering the outages), the probabilities of having a certain number of satellites alive, active (without considering outages) or available (considering outages) per plane and in the constellation, evolution of the age of the satellites per plane and in the constellation, evolution of average number of satellites per status (i.e. nominal, spare, relocating, decommissioned or dead) per plane and in the constellation or average number of corrective launches among others.

3.4. Simulation Settings

Other simulation settings that are specified by the user are the number of samples (constellations) used to compute the probabilities, the simulation time step, and whether the service outages have to be considered in the replenishment or just for post-processing of the simulation outputs.

4. RESULTS OF MODELLING REPLENISHMENT STRATEGIES

As introduced before, the goal of the tool is to produce a replenishment plan able to maintain the service availability over a user-defined threshold. This section will describe some representative, illustrative and significant use cases with the objective of providing the reader with qualitative reference data that give some insight on the capabilities of the Replenishment Planner.

The launches are scheduled whenever a certain figure of merit, which depends on the replenishment strategy selected by the user, violates a user-given condition. Different replenishment strategies can be modelled, and the first three subsections shows the capabilities and typical outputs of the tool through studying three different strategies.

4.1. Preventive Strategy

The preventive replenishment strategy is the most intuitive one. Using this strategy, a launch is scheduled so that the satellite arrives at the target plane just as the service availability violates the user-defined threshold. This means that the launch must be scheduled based on the predicted service availability some time in the future; the date of launch is evaluated as the violation date minus the satellite's transfer time.

The plane with lowest spare capability (i.e. probability of having more than the nominal number of satellites) is chosen as the target plane of the launch. The tool is able to model very complex launch scenarios, with staged injections to multiple planes: in this case the second target plane is defined by the user with its relative position with respect to the primary plane.

The following example corresponds to a constellation of 3 planes, with 14 slots per plane (8 of them nominal, 6 spare) providing a service called "mainService", which undergoes an unscheduled outage of 1 day with a constant failure rate of 0.001 failures per day. Following a failure, the relocation time for a satellite to move from a spare slot to a nominal slot is 10 days. The time step chosen has been 1 day and the number of constellations simulated 150,000.

Figure 1 shows the evolution of the service availability when a preventive replenishment strategy is used. The service availability without outages (active service availability) is used as triggering figure of merit.

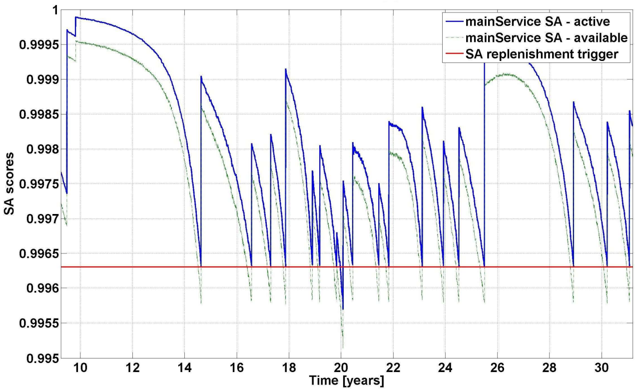

Figure 1: Evolution of service availability over time using a preventive strategy

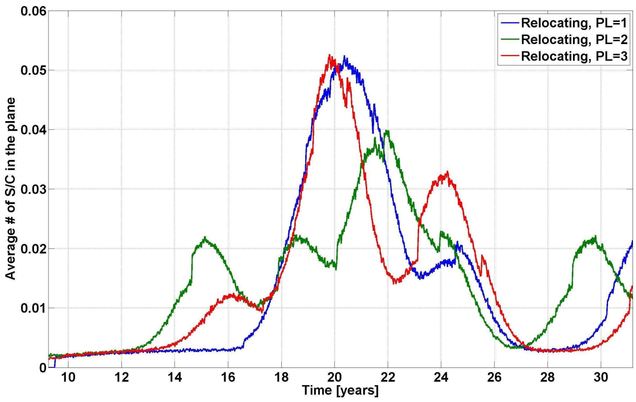

Figure 2: Average number of relocating satellites in each plane

Since a preventive strategy is used, the violations of the threshold are prevented unless a necessary launch cannot be performed due to the presence of some user-defined constraint. In this case, the threshold is violated once, due to the minimum time between launches constraint, which is set to 90 days.

Another graphical output of the Replenishment Planner that provides significant information about of the death rate of the nominal satellites is the average number of satellites in relocation. It can be seen in Figure 2 that during the years of higher replenishment frequency, the number of relocating satellites is higher.

The increment of death rate has an interesting side effect whose impact increases with the number of spare slots in the constellation. For those constellations with a large buffer of spare satellites, the service availability is recovered mostly by means of the relocation instead of by the arrival of new satellites during phases of high death rate. This behaviour is represented in Figure 3, where the evolution of the service availability is shown for two different constellations, one with 2 spare slots per plane and one with 6. In both cases, the relocation time has been set to 25 days.

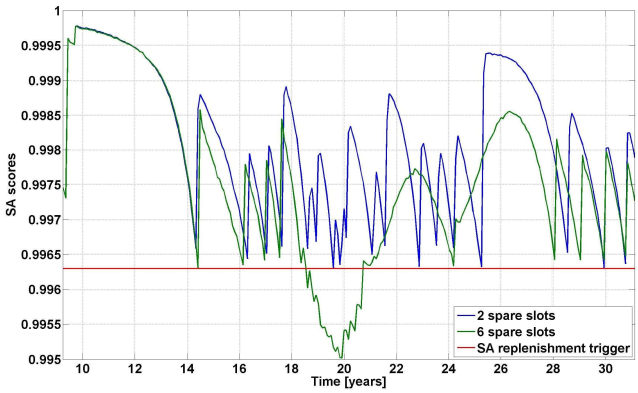

Figure 3: Comparing evolution of service availability between simulations with different spare capacity

The constellation with 2 spare slots is able to keep the service availability over the threshold, whereas the constellation with 6 slots has a threshold violation that lasts around 3 years, where many launches are scheduled (one every 90 days) trying to recover the service availability level. Moreover, the constellation with 2 spare slots schedules 18 launches during the simulation, whereas the other one schedules 19.

This unexpected effect can be explained because in most of the constellation samples relocations are triggered as soon as a nominal satellite dies due to the high number of spare satellites available. The arriving satellites are headed to spare locations instead of to the nominal ones, delaying the service level recovery and muting its response to the arrival of new satellites. In addition, the constellation with 2 spare slots has younger satellites, since the older ones are replaced by the new ones when there are no free slots.

4.2. Preventive-Corrective Strategy

The preventive-corrective strategy is very similar to the preventive one, but it is focused on maintaining the spare capability on each plane, rather than the overall constellation service availability. A launch is triggered when the probability of having more than the nominal number of satellites goes below a user defined threshold.

The following example corresponds to a constellation of 3 planes, with 10 slots per plane, 8 of them nominal, providing a service called “mainService”. As in the previous section, a time step of 1 day, and simulation of 150,000 constellations were used.

In this case, the figure of merit used to trigger launches is the spare capability (probability of having more than 8 satellites alive in a plane) with a threshold of 85%, as shown in Figure 4.

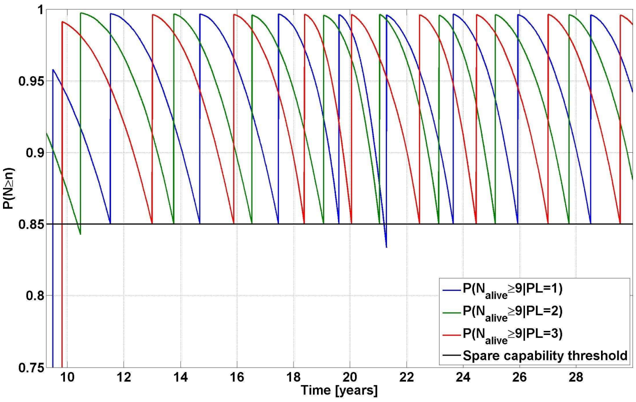

Figure 4: Evolution of per-plane spare capability using a preventive-corrective strategy

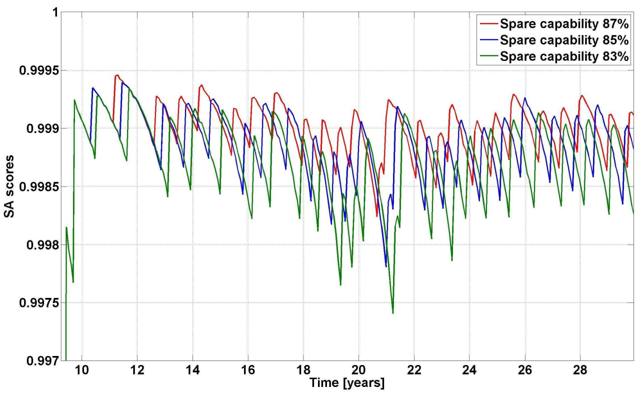

Figure 5: Evolution of service availability over time using preventive-corrective strategy with different thresholds

Figure 5 shows the evolution of the service availability obtained with different values of the spare capability threshold (83%, 85% and 87%). With the threshold set at 85% this strategy is able to maintain the service availability over 99.75%; a higher threshold increases the service availability at a cost of more launches.

Figure 6 shows the percentage of satellites within user-given age intervals, which is very useful to analyse the retirement age of the satellites. It can be seen that, for the considered case, there are very few satellites between 15 and 20 years of age.

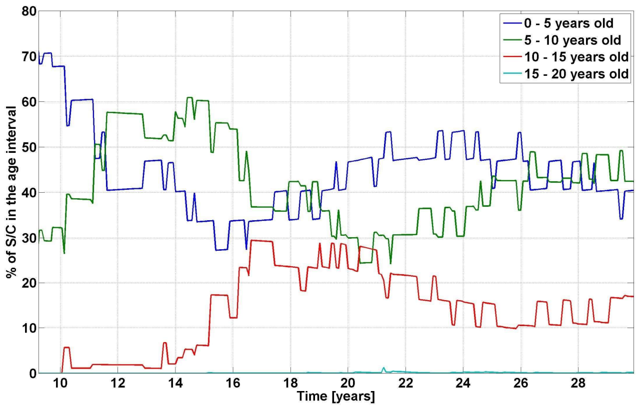

Figure 6: Satellite age distribution

4.3. Corrective Strategy

The corrective strategy presents two major differences with respect to the preventive and preventive-corrective ones. The first one is that it schedules a launch whenever one of the constellations composing the statistical sample loses its spare satellites in any of its planes, and incidentally there are not on-going launches targeted to replenish that plane. Thus, the launches are not applied to all the constellations composing the statistical sample (as it happened in the previous strategies) with the consequence that the simulation output launches plan is given as an averaged replenishment plan on the statistical sample.

In addition, the launches are scheduled ahead in time (not backwards): when the last spare satellite of a plane is lost, the corrective launch is requested with a delay between the launch request (violation of the triggering condition) and the launch date (and therefore the arrival of the satellite in its in-orbit slot).

4.4. Combined Strategies

The Replenishment Planner has been designed and implemented in order to allow robust combinations of preventive and corrective, and preventive and preventive-corrective strategies.

4.5. Sensitivity Analysis

The Replenishment Planner can be used to perform sensitivity analysis over different parameters of the simulation. For example, the influence of a parameter can be analysed using a reference replenishment plan. Figure 7 shows the service availability evolution obtained with two different satellite reliability curves, but using the same replenishment plan. The blue curve corresponds to a simulation performed assuming higher reliability satellites, so the level of service availability provided is higher; whereas the green one has been obtained using satellites with lower reliability.

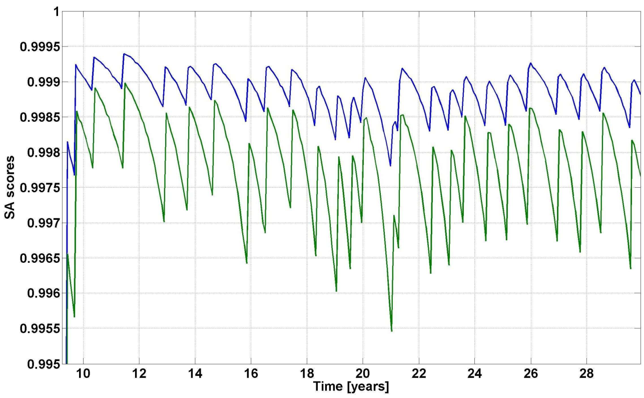

Figure 7: Service availability evolution with different reliability curves

4.6. Predictor-Corrector Replenishment Planner

A Monte Carlo analysis has been performed launching 250 simulations of the Predictor-Corrector Replenishment Planner in sensitivity mode using the first scenario of section 4.1.

The histogram of Figure 8 shows the number of launches performed each month in total over all 250 Monte Carlo cases. All the first launches of each simulation are represented in dark blue, the second launches in a lighter blue, etc. Almost all the simulations (94.8%) need at least 16 launches to keep the service availability, very few (4.4%) need more than 17. It is also relevant to note that the distribution of first launch dates is very narrow: almost half of them occur in the same month. Conversely, the later ones are more spread out.

In addition, the evolution of the service availability computed as an average among the constellation of each simulation is a valuable output, as shown in Figure 9. The violations of the threshold by the active service availability occur during 2032; the longest violation (that corresponds to the sharp decay of available service availability) lasts for approximately 20 days. The fact of observing very few violations of the threshold is typical of a sensitivity analysis.

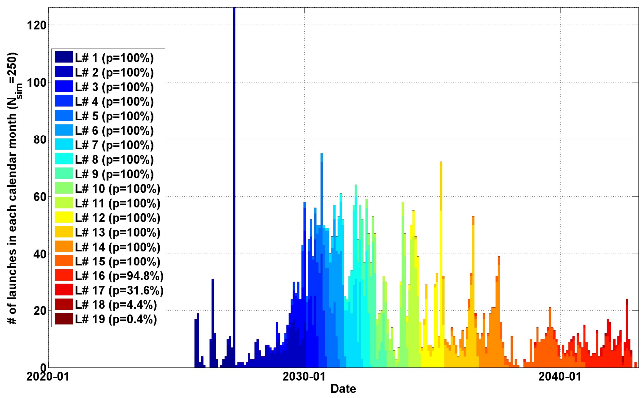

Figure 8: Monthly histogram of the launches

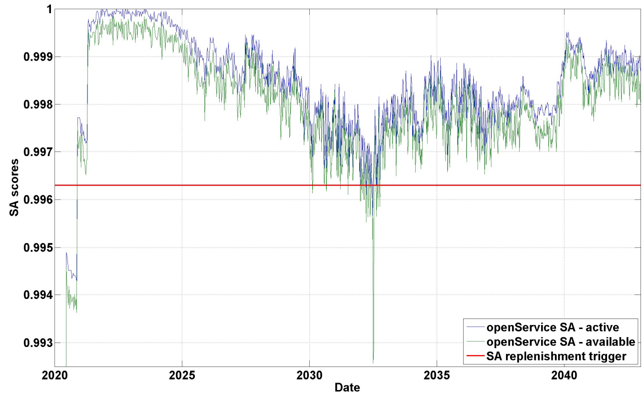

Figure 9: Service availability evolution extracted from the Monte Carlo simulations

5. PERFORMANCE

Replenishment Planner is an extremely demanding tool in terms of computational work-load. The analytical approach to the computation of the constellation state probabilities is practically unfeasible with the state-of-the-art technology. A statistical approximation to the problem, even if feasible, can require hundreds of thousands of constellation samples in order to have the required level of confidence of the estimated probabilities, depending on the complexity of the problem to be studied. Therefore, the Replenishment Planner has been designed with a variable time-step that is able to provide excellent runtime performances with a very low impact on the evaluation of the solution.

The parameters that have most impact on the computational time are the duration of the simulation, the time step and the number of constellation samples. Figure 10 provides an example of the speed up that can be achieved, showing the same calculation as Figure 4, but with the simulation run using a much larger time step (30 days instead of 1 day). It can be seen that the spare capability curves are not so close to the threshold. However, if the evolution is compared with the results obtained with 1-day time step (Figure 4), the curves have a very similar pattern thanks to the variable time step implemented. The improvement in runtime performances is directly proportional to the increase of the time step, therefore it is possible to have a 30-fold speed-up factor by increasing the time step length without qualitatively affecting the simulation results.

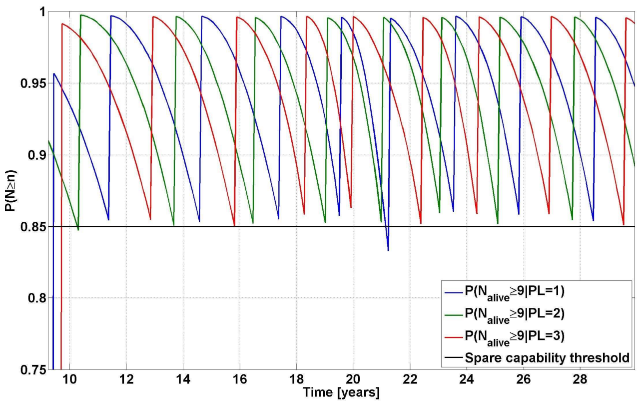

Figure 10: Evolution of per-plane spare capability using a time step of 30 days

A typical low fidelity simulation of 30 years duration, 30 days time step and 150,000 constellation samples requires around 10 minutes, whereas the same simulation performed with a high fidelity 1 day time step would require approximately 4.5 hours.

The simulator is also able to distribute the computational processes to the available CPU cores when batch, parametric and/or Monte Carlo analyses are performed. Thanks to the task distribution capabilities of openSF, the computational time can be drastically reduced when several analysis have to be performed.

6. APPLICATION TO OTHER CONSTELLATION TYPES

Service availability is an abstract geometrical concept that measures the geographical coverage provided by the constellation payloads considered. Therefore, it can be applied not only to GNSS, but to any constellation focused on an area of interest. For instance, communications mega-constellations are a suitable candidate to be analysed with the Replenishment Planner once the service availability coefficients are obtained by computing the coverage of the area of interest with different in-orbit configurations. Moreover, it would be possible to use the tool for Earth Observation constellations equipped with passive or active sensors; also in this case, the service availability coefficients should be evaluated according to the in-orbit configuration defining the constellation and the sensors coverage.

A further possibility provided by the high modularity of the tool would be to replace the service availability with a different control variable that better measures a constellation's performance. For example an LEO Earth Observation constellation might measure the value of service delivered considering revisit frequency for points of interest, rather than a binary measure of availability.

7. FUTURE DEVELOPMENTS

Most of the future developments of the Replenishment Planner are focused on improving the fidelity of the models used to compute the failure time of the spacecrafts. For instance, the computation of the reliability of the services can be improved by splitting down each service in components (each with its own reliability and outages definition) and implement a reliability chain. This way, it would be possible to consider failures of critical components that affect to multiple services.

In addition, a brand new predictive strategy could be implemented, based on the estimation of the expected lifetime of the spacecrafts. This new replenishment strategy would trigger launches independently to each constellation in the simulation whenever the expected number of satellites is not enough to guarantee a proper service level. The computation of the expected lifetime of the satellites might vary from very simple implementations to other more complex, that would take advantage of the service split down in components, based on the monitoring of the status of the redundant critical components. Furthermore, the combination of the predictive strategy and the corrective one, that would mitigate the prediction failures, seems quite promising.

8. CONCLUSION

The Replenishment Planner is a highly flexible simulator able to provide valuable statistical information about the replenishment needs of a constellation and the quality of the service provided: a sequence of replenishment launches that guarantees the availability of the designed service above a user-defined threshold. It is highly configurable and able to model a wide variety of constellations in terms of in-orbit geometry, spacecraft and services models, operational launch characteristics and constraints.

In addition, the tool can analyse the goodness of different launch scenarios and perform sensitivity and robustness analyses on multiple input parameters to evaluate its impact on the resulting replenishment plan and the service availability.

The Replenishment Planner has been conceived and implemented taking into account operational constraints, thus is able to perform complex and reliable analyses very quickly, nevertheless providing the possibility to have more accurate results within reasonable time spans.

The Replenishment Planner represents a powerful tool to support system analyses in the frame of constellations maintenance and a forerunner simulator for an operational tool to be possibly used as support for day-to-day on-ground operations whenever a replenishment strategy for constellation maintenance would be required.

9. REFERENCES

[1] OpenSF, European Space Agency, https://eop-cfi.esa.int/index.php/opensf